

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE  
UTILITY PATENT APPLICATION TRANSMITTAL UNDER 37 CFR 1.53(b)

## Address to:

Assistant Commissioner for Patents

Box Patent Application

Washington, DC 20231

Attorney Docket No.

2207/7942

Inventor(s)

MARR et al.

Express Mail Label No.

EL502647119US

Total Pages

21

JC678 U.S. PTO

09/490172



## Title of Application:

## ESTABLISHING PRIORITY IN A PROCESSOR OR THE LIKE

## Transmitted with the patent application are the following:

- 2 Page(s) Transmittal form (and one copy)  
15 Page(s) Cover Page (1), Specification (8), claims (4), abstract (1)  
4 Page(s) Informal drawing  
Page(s) Declaration and Power of Attorney (unsigned)  
Page(s) Recordation of Assignment and Assignment form  
Page(s) Information Disclosure Statement (IDS) (copies of citations not included in number of pages)  
Page(s) Certified copy of:

This application is a Continuation / Continuation-in-Part / Divisional of prior application Serial No.         , filed         .

## Fee calculation for large entity:

	No. Filed	No. Allowed	No. Extra	Rate	Fee
Basic Fee					\$690.00
Total Claims	20	20		× 18.00	
Independent Claims	4	3	1	× 78.00	\$78.00
Multiple Dependent Claim				+ 260.00	
				Assignment	
				Total	\$768.00

The Commissioner is hereby authorized to charge payment of any fees associated with this communication or arising during the pendency of this application per 37 CFR §1.16–1.21, or to credit any overpayments, to Deposit Account 11-0600.

## Express Mail Certificate

I hereby certify that the above paper/fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated below and is addressed to the Assistant Commissioner for Patents, Washington, DC 20231

Date of deposit: January 22, 2000

Person mailing paper/fee: Shawn O'Dowd

Signature

## Respectfully submitted,

Shawn W. O'Dowd (Reg. No. 34,687)  
Attorney for Applicant(s)

Kenyon & Kenyon  
333 W. San Carlos Street, Suite 600  
San Jose, CA 95110  
(408) 975-7500 - Telephone  
(408) 975-7501 - Facsimile

U.S. Patent Application

Title: Establishing Thread Priority in a Processor  
or the Like

Inventors: Debbie Marr et al.

## **Establishing Thread Priority in a Processor or the Like**

### Background of the Invention

The present invention pertains to the operation of a processor or the like. More particularly, the present invention pertains to establishing priority of a thread in a multi-threaded processor.

As is known in the art, a processor includes a variety of sub-modules, each adapted to carry out specific tasks. In one known processor, these sub-modules include the following: an instruction cache, an instruction fetch unit for fetching appropriate instructions from the instruction cache; decode logic that decodes the instruction into a final or intermediate format, microoperation logic that converts intermediate instructions into a final format for execution; and an execution unit that executes final format instructions (either from the decode logic in some examples or from the microoperation logic in others).

Programming code to be executed by the processor can sometimes be broken down into smaller components referred to as "threads." A thread is a series of instructions whose execution achieves a given task. For example, in a video phone application, the processor may be called upon to execute code to handle video image data as well as audio data. There may be separate

code sequences whose execution is designed to handle each of these data types. Thus, a first thread may include instructions for video image data processing and a second thread may be instructions for audio data processing.

In some multi-threaded processors, the processor may switch between execution of two or more threads. In other multi-threaded processors, the threads may be executed simultaneously. In either of these processors, there is no delineation between how the threads are treated. In particular, code from one thread is given the same priority as code from another thread. This could lead to a negative impact on overall system performance, especially when execution of critical code is suspended or slowed by the execution of non-critical code.

In view of the above, there is a need to establish priority between two or more threads.

### Summary of the Invention

This and other needs are satisfied by embodiments of the present invention. In one embodiment, a method of establishing thread priority in a processor is presented where a value in memory is assigned to indicate which of a plurality of threads has a higher priority.

### Brief Description of the Drawings

Fig. 1 is a block diagram of a computer system operated according to an embodiment of the present invention.

Fig. 2 is a block diagram of a portion of a processor system constructed according to an embodiment of the present invention.

Fig. 3 is a block diagram of a portion of a processor system constructed according to an embodiment of the present invention.

Fig. 4 is a block diagram of a bus system constructed according to an embodiment of the present invention.

#### Detailed Description

Referring to Fig. 1 a block diagram of a computer system operated according to an embodiment of the present invention is shown. In this example the computer system 1 includes a processor 3 which is capable of executing code stored in memory 5. In this example, memory 5 stores code for several threads, such as code for thread 0 (8), thread 1 (9), etc. As known in the art, code for two threads may be part of user applications and for the operating system.

Referring to Fig. 2, a block diagram of a processor system (e.g., a microprocessor, a digital signal processor, or the like) operated according to an embodiment of the present invention is shown. In this embodiment, the processor is a multi-threaded processor where the execution unit 41 is theoretically divided into two or more logical processors. As used herein, the term "thread" refers to an instruction code sequence. For example, in a video phone application, the processor may be called upon to execute code to handle video image data as well as audio data. There may be separate code sequences whose execution is designed to handle each of these data types. Thus, a first thread may include instructions for video image data processing and a second thread may be instructions for audio data processing. In this example, there are one or more execution units (e.g., including execution unit 41), which may execute one or more instructions at a time. The processor system 10, however, may be treated as two logical processors, a first logical processor executing instructions from the first thread and a second logical processor executing instructions from the second thread.

In this embodiment of the processor system 10, instructions and/or bytes of data are



predicted, the appropriate instructions are immediately available for the execution unit. If the branch prediction is incorrect, then the instructions in the pipeline are incorrect and must be flushed out and the appropriate instructions loaded into the pipeline.

According to an embodiment of the present invention, an operating system that supports multiple thread execution may set one or more bit flags in memory (e.g., memory 4 provided in processor 3) to indicate that a particular thread is to be given priority over another. For example the lower four bits of the APIC (Advanced Programmable Interrupt Controller) TPR (Task Priority Register) Register for each thread may be used to set thread priority. The operating system may access the APIC TPR when it schedules a task for processing. The upper four bits of the APIC TPR register are used by the operating system to set priority as between interrupts. The setting of the lower four bits in the APIC TPR register serves as a hint to the processor system that one or more threads are to be given a higher priority in using the processor resource. Determining which thread is to be given priority is application specific. As an example, in a video conferencing application including code for processing video and audio data, the operating system may assign a higher priority to threads from either the video or audio data processing code rather than threads from the other data processing code. In this embodiment of the present invention the variable or flag Thread0Priority (e.g., the lower four bits of thread 0's APIC TPR register) in memory 4 is set to a higher value if thread 0 is to be given a higher priority than other threads and variable or flag Thread1Priority (e.g., the lower four bits of thread 1's APIC TPR register) in memory 4 is set to a higher value if thread 1 is to be given a higher priority than other threads. The four bits will be set to the same value if the threads are to be given the same priority.

According to an embodiment of the present invention, the setting of the priority bits or

flags allows processing of one thread to take precedence over other threads of lower priority value. Threads of the same value are given the same priority. As a first embodiment of the present invention, control logic 17, 26, 33, 40, and 45 may be used to select which thread is given access to the next stage of processing. For example, control logic 17 may be used to select how many instructions from each thread are to be forwarded to decode unit 21 (based on the values stored in flags Thread0Priority and Thread1Priority). The length of the instruction, in bytes, depends on what type of processor is being used. For a RISC (Reduced Instruction Set Code) processor, each instruction is typically one byte. For a CISC (Complex Instruction Set Code), the instruction may be one to fifteen bytes or longer (e.g., for example the IA-32 architecture processors of Intel Corporation, the maximum instruction length is 15 bytes long). For a CISC instruction, the decode unit 21 may decode the instruction into a number of microinstructions (e.g., a one byte CISC instruction may be decoded into three microinstructions).

In this example, two threads are shown, though the invention is not so limited. If the threads are to be treated without assigning a priority to either one, then one method for handling instructions from each thread is to retrieve a predetermined number, X, of instruction bytes from a first thread then continue to retrieve bytes from the first thread until a taken branch in the code (i.e., based on a branch prediction unit) or some other set number of instruction bytes, whichever comes first. When reaching this point for the first thread, then the control logic switches to the other thread for processing in a similar manner.

According to this embodiment of the present invention, the predetermined value may be set for each thread so as to give priority of processing to one thread over the other. Referring to Fig. 3, a free-running timer or down-counter 68 is provided that is loaded, upon a thread switch,



with either a first, low value 69 (e.g., 15), or a second, high value 70 (e.g., 30). Thus, if thread 0 is given a higher priority than thread 1, then when switching to the fetch of thread 0 instructions from queue 13, the high value is loaded into the counter 68, and when switching to the fetch of thread 1 instructions, the low value is loaded into the counter 68. Once counter 68 reaches zero, instructions are loaded into decode unit 21 from the current thread until a taken branch is reached or a preset number of additional instructions is reached; control logic 17 then switches to the next thread. The values for the high and low registers 69, 70 are freely programmable in this embodiment (e.g., through control registers). Referring to Fig. 2, if thread 0 is assigned a higher priority and instructions from this thread are being loaded from a source other than queue 13, then the priority feature shown in Fig. 3 may be disabled so that instructions from the lower priority thread can be fetched and decoded without interruption if desired. Also, if no instructions or bytes are available for a high priority thread, then instructions or bytes could be loaded from the lower priority thread so as to not degrade efficiency of the processor system. In other words, threads of lower priority may be given greater access to a resource compared to a thread with a higher priority when the higher priority thread is not using the resource.

As stated above, each control logic may be used to control which thread will be given access to a particular unit or resource in the processor system. The use of control logic 17 to control which thread supplies instructions or bytes to decode unit 21 is one example of many for this embodiment of the present invention. For instance, control logic 33 may be used to control the number N of instructions that are provided from each thread from queue 29 of the trace cache/MSROM unit to the rename allocation unit 35.

The present invention can be extended to other aspects of a computer system. For example, access to a cache memory may be controlled so that a higher priority thread is given

greater access to the cache (e.g., by assigning more ways in the cache to the higher priority thread). If there is a resource that includes six buffers, then four of these buffers may be assigned to the high priority thread and two assigned to the low priority thread to increase performance of the high priority thread. An example of assigning priority access to a cache is shown in pending application number 09/224,377 filed on December 31, 1998.

Also, to the extent thread execution requires use of a bus or results in use of a bus, the indication of thread priority may be used to provide more access to the higher priority thread. Thus, in this case the shared resource between thread execution is a bus. For example, referring to Fig. 4, a processing system includes a memory (L1 cache) 83 coupled to an instruction fetch unit 84. The instruction fetch unit provides instructions or bytes of data to decode unit 85, which in turn is coupled to execution unit 87. Execution unit 87 is coupled to memory 83. Memory 83 is further coupled to bus 81 via a bus unit 82. Bus unit 82 includes a thread 0 request queue and a thread 1 request queue. According to an embodiment of the present invention, control logic in bus unit 82 controls which bus requests by threads 0 and 1 are transacted on bus 81 (e.g., a system bus). For example, the control logic may alternate between five consecutive accesses to thread 0 and two consecutive accesses to thread 1 when thread 0 has a higher priority (e.g., by setting the appropriate flag in memory as described above). The number of bus access assigned to high and low priority threads may be freely programmable by the user-application, for example.

Using thread priority as described above allows an increase in performance for the execution of varied applications. Assigning a higher priority to code that requires faster, timely execution, results in a better distribution of the processor resource to the applications using it.

Although several embodiments are specifically illustrated and described herein, it will be



What is claimed is:

1. A method of establishing thread priority in a processor comprising:  
assigning a value in memory to indicate which of a plurality of threads has a higher priority.
2. The method of claim 1 further comprising:  
allocating a resource between said plurality of threads depending on a priority assigned to each thread.
3. The method of claim 2 wherein in said allocating step, a first thread is given greater access to a resource than other threads when said first thread is assigned a higher priority than said other threads.
4. The method of claim 2, wherein in said allocating step, the other threads are given greater access to the resource than the first thread when said first thread is assigned a higher priority than the other threads and is not using said resource.
5. The method of claim 3 wherein said resource is a unit in a processor system.
6. The method of claim 5 wherein said resource is a decode unit.
7. The method of claim 6 further comprising:  
providing instructions from a first thread to a first queue;

providing instructions from a second thread to a second queue;  
supplying a first number of instructions to said decode unit from said first queue;  
supplying a second number of instructions to said decode unit from said second queue;  
selecting said first and second numbers based on said value in memory.

8. The method of claim 3 wherein said resource is a bus.

9. The method of claim 8 further comprising:

providing bus requests from a first thread to a first queue;  
providing bus requests from a second thread to a second queue;  
servicing a first number of bus requests from the first queue;  
servicing a second number of bus requests from said second queue; and  
selecting said first and second numbers based on said value in memory.

10. A method of establishing thread priority in a processor comprising:  
assigning a value in an APIC TPR register for a thread via execution of operating system code to indicate which of a plurality of threads has a higher priority.

11. An apparatus for establishing thread priority in a processor comprising:  
a memory to store a value to indicate which of a plurality of threads has a higher priority.

12. The apparatus of claim 11 further comprising:  
a resource allocated between said plurality of threads depending on a priority assigned to

each thread in said memory.

13. The apparatus of claim 12 wherein a first thread is given greater access to a resource than other threads when said first thread is assigned a higher priority than said other threads.
14. The apparatus of claim 12 wherein the other threads are given greater access to the resource than the first thread when said first thread is assigned a higher priority than the other threads and is not using said resource.
15. The apparatus of claim 13 wherein said resource is a unit in a processor system.
16. The apparatus of claim 15 wherein said resource is a decode unit.
17. The apparatus of claim 16 further comprising:
  - a first queue to store instructions from a first thread;
  - a second queue to store instructions from a second thread;
  - control logic coupled to said first and second queues and said decode unit, said control logic to permit a first number of instructions to be transferred to said decode unit then a second number of instructions to be transferred to said decode unit, said first and second numbers being selected based on said value in memory.
18. The apparatus of claim 13 wherein said resource is a bus.

19. The apparatus of claim 18 further comprising:

a bus unit including

a first queue storing bus requests from a first thread;

a second queue storing bus requests from a second thread;

control logic coupled to said first and second queues, said control logic to control servicing of a first number of bus requests from the first queue and a second number of bus requests from said second queue, said first and second number being selected based on said value in memory.

20. An apparatus for establishing thread priority in a processor comprising:

an APIC TPR register for a thread wherein execution of operating system code causes a value to be stored in said register to indicate which of a plurality of threads has a higher priority.

### Abstract of the Disclosure

In a multi-threaded processor, one or more variables are set up in memory (e.g., a register) to indicate which of a plurality of executable threads has a higher priority. Once the variable is set, several embodiments are presented for granting higher priority processing to the designated thread. For example, more instructions from the higher priority thread may be executed as compared to the lower priority thread. Also, a higher priority thread may be given comparatively more access to a given resource, such as memory or a bus.



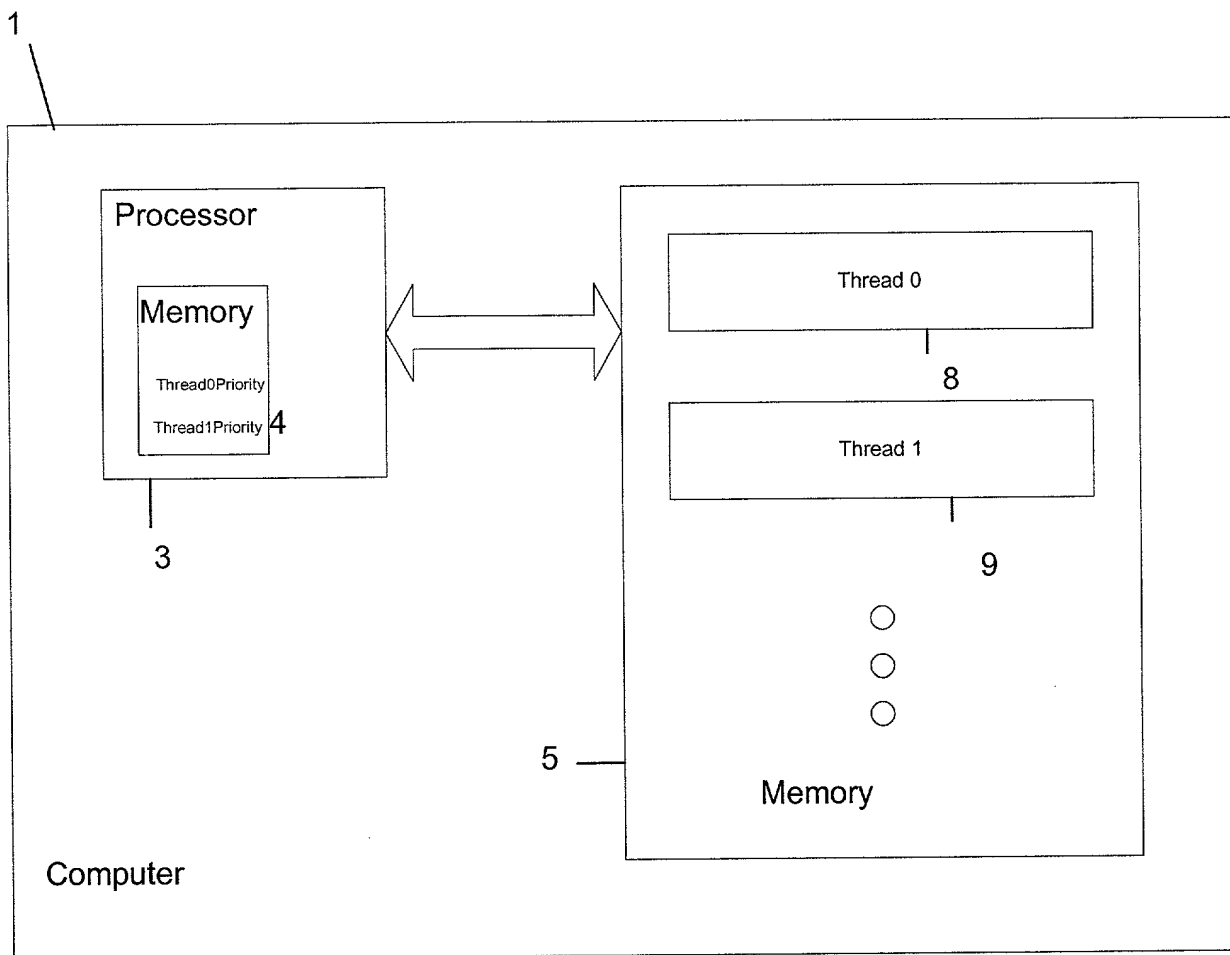


Fig. 1




Fig. 2




Fig. 3




Fig. 4